

# Valuation of Energy Storage: An Optimal Switching Approach

René Carmona

Department of Operations Research and Financial Engineering, also with Bendheim Center for Finance,  
Princeton University, Princeton, NJ 08544 rcarmona@princeton.edu,

Mike Ludkovski

Department of Statistics and Applied Probability, University of California, Santa Barbara, CA 93106  
ludkovski@pstat.ucsb.edu, <http://www.pstat.ucsb.edu/faculty/ludkovski>

We consider the valuation of energy storage facilities within the framework of stochastic control. Our two main examples are natural gas dome storage and hydroelectric pumped storage. Focusing on the timing flexibility aspect of the problem we construct an optimal switching model with inventory. Thus, the manager has a constrained compound American option on the inter-temporal spread of the commodity prices. Extending the methodology from Carmona and Ludkovski (2008), we then construct a robust numerical scheme based on Monte Carlo regressions. Our simulation method can handle a generic Markovian price model and easily incorporates many operational features and constraints. To overcome the main challenge of the path-dependent storage levels two numerical approaches are proposed. The resulting scheme is compared to the traditional quasi-variational framework and illustrated with several concrete examples. We also consider related problems of interest, such as supply guarantees and mines management.

*Key words:* gas storage; optimal switching; least squares Monte Carlo; hydro pumped storage; impulse control, commodity derivatives

*History:* First Version: July 2005; This Version: August 24, 2008

---

## Acknowledgments

We thank the participants of the Banff BIRS Workshop 07w-5502 “Mathematics and the Environment” and Zhenwei J. Qin for many useful comments and discussions. We also thank the anonymous referees whose feedback led to much improved presentation.

## 1. Introduction.

While classical financial contracts such as stocks and bonds are paper assets, ownership of commodities entails physical storage. As a result, the modern commodities industry incorporates an extensive storage infrastructure, including natural gas salt domes, liquified natural gas (LNG) storage tanks, precious metal repositories and hydroelectric reservoirs. In the last decade, with the ongoing deregulation of these industries, storage facilities have also acquired an important role

in the commodity financial markets. Storage allows for inter-temporal transfer of the commodity and permits exploitation of the fluctuating market prices. The basic principle is to ‘buy low’ and ‘sell high’, such that the realized profit covers the intermediate storage and operating costs. Since the profitable opportunities are driven by price volatility, the storage facility grants its owner a calendar *straddle* option.

Traditionally, storage facilities have been owned by major players in the respective industries who had the enormous capital typically needed to build and maintain them. However, with the liberalized markets, all participants have nowadays the opportunity to *rent* a storage facility with an eye towards speculation on prices and aggressive profit maximization. For example, with respect to natural gas de Jong and Walet (2003) write that “natural gas storage is unbundled, ... [and] offered as a distinct, separately charged service. ... Buyers and sellers of natural gas have the possibility to use storage capacity to take advantages of the volatility in prices”.

The aforementioned price volatility can be either systemic or speculative. For instance, the natural gas market exhibits strong seasonality since the main consumer group is households that use gas for winter heating. Thus, natural gas demand (and prices) has a systematic spike in the cold season, often on the order of 30%-50%. In contrast, in say silver, the volatility in prices is almost entirely speculative, but nevertheless can sometimes lead to price swings of 100% within one year, cf. the Hunt brothers episode in late 1970s (Pirrong 1996).

The seasonality effects present in certain markets lead to *intrinsic* value of storage that can be locked-in by a static purchase and sale of forward contracts. For instance, a typical July-January forward spread in natural gas is on the order of \$1/MMBtu and can be easily realized by a simple one-time transaction. However, the presence of an increasingly liquid short-term market permits further *dynamic* optimization. Intuitively, the manager holds timing options that allow her to optimally exploit opportunities that appear as market prices evolve. Capturing this optionality is crucial in the present competitive markets, especially with the entry of non-energy players who rent facilities with the sole goal of maximizing profit (as opposed to old-fashioned participants who also have strategic aims). Moreover, with the growing importance of energy commodities, sophisticated valuation of energy storage becomes an integral aspect of functioning financial markets.<sup>1</sup>

Thus, it is necessary to be able to compute the financial *extrinsic* value of such flexibility. Namely, how much should one pay to gain control of a storage facility for a period of  $T$  years? The simple question above hides the associated modeling difficulties. Indeed, the owner faces a multitude of

<sup>1</sup> For instance, the \$5 billion loss reported by Amaranth LLC. in Fall 2006 was due to a poorly managed bet on the March-April calendar spread in natural gas, which is in turn driven by actions of storage managers during the winter.

optionalities and constraints that interact in a nonlinear fashion. First, the purchases and sales can be done immediately, using spot prices, or forward in time using forward prices. Second, the bought commodity must be put on the inventory. As a result, inventory capacity limits, as well as storage costs, delivery charges and other operational and engineering constraints become crucial. However, the latter are intrinsically path-dependent in terms of the storage strategy adopted. Finally, the manager may be exposed to margin requirements (if the commodity is bought with credit), mechanical break-downs and other external events.

To overcome these challenges the current literature on commodity storage has largely proceeded in two different directions. The basic practitioner methods have been based on the traditional option-pricing approach. Thus, one makes (often drastic) simplifications to shoehorn the problem into the option pricing framework. For instance, gas storage can be reduced to a collection of calendar Call options paying out the spread between gas prices today and  $k\Delta t$  years from now, with  $k = 1, \dots, T/\Delta t$  (Eydeland and Wolyniec 2003). After such a reduction the extensive existing machinery of derivative pricing can be imported. One gains intuition and computational speed but ignores key operational constraints (such as dynamic capacity limits), since the calendar Calls are priced independently of each other. Furthermore, the method is *ad hoc*, requiring heuristic adjustments to correct for model assumptions. Alternatively, various stochastic programming algorithms (Nowak and Römisich 2000, Doege et al. 2006) have been considered, especially for hydrothermal systems. These methods maintain the flexibility of incorporating realistic constraints, but instead discretize the set of future scenarios. While powerful, stochastic programming suffers from non-scalability with respect to number of scenarios and time-steps used.

To properly account for the interdependence between the timing optionality of the manager in choosing the purchase and sale times and the inventory constraints, one must consider the full stochastic control framework. This leads to a Bellman dynamic programming equation for the value function. From here one may apply the Hamilton-Jacobi-Bellman theory, translating the problem into a quasi-variational partial differential equation (pde) formulation. Such models were studied in Ahn et al. (2002) and Thompson et al. (2003) and solution is then obtained via standard numerical solvers. However, the path-dependency due to presence of inventory implies that the pde is degenerate (convection-dominated) and therefore extra care is necessary. Moreover, the implementation is necessarily price-model dependent and consequently not robust. One approach that mitigates these challenges was recently proposed in Chen and Forsyth (2007).

In this paper, we also adopt the stochastic control formulation. However, in contrast to the pde methods above, we proceed to a probabilistic solution based on optimal multiple stopping

problems. This perspective allows us to obtain an efficient *simulation-based* numerical method for valuing energy storage on a finite horizon. The method is flexible and not tied to a particular class of asset prices; in fact we abstract from asset dynamics and take as exogenous the (multi-dimensional) Markov price process for the commodity. Thus, use of complex price dynamics, such as jump-diffusions, several factors, etc. has only a marginal impact on the efficiency of the algorithm.

Compared to previous approaches, our method has several advantages. First, we maintain rigorous modeling of the operational constraints while considering the entire set of future scenarios of commodity prices. Moreover, in contrast to pde solvers which suffer from the curse of dimensionality, our scheme can easily handle multi-dimensional settings. In terms of performance, our scheme is competitive with the pde solvers in one-dimension and is clearly superior in higher dimensions (which are essential in a realistic model, see Section 7). Thanks to its scalability, the algorithm is easily extendable and therefore suitable for realistic use. Thus, our main contribution is a robust numerical method that remains on firm theoretical grounds of stochastic control while bridging the gap to practitioner needs.

To be concrete, from now on we focus on the representative example of controlling a natural gas salt dome facility; other applications are addressed in Sections 6 and 7. The rest of the paper is structured as follows. Section 2 describes the stochastic control model we use and its relation to existing literature. Section 3 summarizes the theoretical solution method which is then implemented in Section 4. After outlining the numerical scheme, we proceed to illustrative examples in Section 5. Sections 6 and 7 discuss hydroelectric pumped storage and several problems in natural resource management and demonstrate that our methodology is applicable to a wide variety of real options encountered among commodity derivatives. Finally, Section 8 concludes and outlines future projects.

## 2. Stochastic Model.

Natural gas storage is currently the most widespread class of commodity storage infrastructure in the US<sup>2</sup> (FERC 2004). A variety of storage options, including depleted gas fields, aquifers, salt domes and artificial caverns are available. In 2006 over 400 such facilities existed in the US and a substantial portion are contracted out for periods of 6–60 months. In the near future, the industry will expand even more with the rolling out of LNG technology and associated storage in North America (see Geman (2005) for the general trends and organization of the gas universe.). In this article we will specifically focus on the case of salt domes which permit the highest rates of injection and withdrawal and therefore contain the most timing optionality (see Table 3).

<sup>2</sup> Throughout we focus on the North American markets and use imperial system units.

A salt dome is an underground natural cave that can store several billion cubic feet of gas (Bcf). It is connected via pumps to the national pipeline system which allows to inject/withdraw gas at a deliverability rate of  $0.1 - 0.4Bcf$  per day. Taking the point of view of the renter, or *manager* of such a cave, we now wish to maximize economic value by optimizing the dispatching policy, i.e. dynamically deciding when gas is injected and withdrawn, as time and market conditions evolve. We assume that the manager is rational and risk-neutral and aims to maximize total expected revenue over the finite horizon of her rental. We moreover assume that the respective financial markets are liquid and the manager is a price-taker (the situation of price impact is treated in Section 7).

The ingredients of our model can be now listed as:

- Time horizon  $T$ , with a stipulation for the final state of the facility, see (6).
- Market gas prices given by a Markov continuous-time stochastic process  $(G_t)$ ,  $G_t \in \mathbb{R}^d$ , quoted in dollars per million of British thermal units (MMBtu), with  $1 \text{ Bcf} \equiv 10^6 \text{ MMBtu}$ .
- Level of inventory in storage denoted by  $C_t$ .
- Finite cave capacity represented by  $c_{min} \leq C_t \leq c_{max}$ .
- Constant discount (interest) rate  $r$ .
- Three possible operating regimes of the storage facility: injection, storage and withdrawal.
- Denote by  $a_{in}(C_t)$  the injection rate, quoted in Bcf per day. Injection of  $a_{in}(C_t)$  Bcf of gas, requires the purchase of  $b_{in}(C_t) \geq a_{in}(C_t)$  Bcf on the open market.
- Similarly the withdrawal rate is labelled  $a_{out}(C_t)$  and causes a market sale of  $b_{out}(C_t) \leq a_{out}(C_t)$  Bcf.
- Capacity charges  $K_i(t, C_t)$  in each regime that represent direct storage costs, delivery charges, various O&M costs and seepage losses.

The case  $b_i \neq a_i$  indicates gas loss during injection/withdrawal (typically on the scale of 0.25% – 1% for salt dome storage). The transmission rates  $a_i, b_i$  themselves are fixed by the physical characteristics of the facility; they are a function of  $C_t$  and are based on gas pressure laws (Thompson et al. 2003).

REMARK 1. Typically,  $G_t$  would represent the price at time  $t$  of the near-month forward contract, which is by far the most liquid contract on the market<sup>3</sup>. However, given a variety of quoted gas prices (spot, balance-of-the-month, futures, etc.), we remain agnostic about the precise interpretation of the  $(G_t)$  process. The driving process  $(G_t)$  may also include longer maturity forwards.

<sup>3</sup> Recent daily volume on NYMEX has been over 90,000 contracts, with more than 50% of the trades in the near-month.

Unfortunately, forward selling is problematic, since the sale price is locked-in in advance, while the inventory only changes at delivery time. We assume for simplicity that any purchase or sale is immediately reflected in the current inventory.

Label the three regimes above as  $i \in \{-1, 0, 1\}$  and denote by  $\psi_i(G_t, C_t)$  the payoff rate (in \$/year) from running the facility in regime  $i$ . Then  $\psi_i$ 's and the corresponding volumetric changes in inventory are given by

$$\begin{cases} \text{Inject: } \psi_{-1}(t, G_t, C_t) = -G_t \cdot b_{in} - K_{-1}(C_t), & dC_t = a_{in}(C_t) dt, \\ \text{Store: } \psi_0(t, G_t, C_t) = -K_0(C_t), & dC_t = a_0(C_t) dt, \\ \text{Withdraw: } \psi_1(t, G_t, C_t) = +G_t \cdot b_{out} - K_1(C_t), & dC_t = -a_{out}(C_t) dt. \end{cases} \quad (1)$$

In principle, the facility can also be operated at a sub-maximal transfer rate, however when the monetary reward is linear in the pumping rate as in (1), it is always optimal to inject/withdraw at maximum speed. This is the so-called ‘bang-bang’ property of stochastic control problems (Øksendal and Sulem 2005).

Many possibilities exist for the form of  $(G_t)$  and there is much recent debate (see Eydeland and Wolyniec (2003), Cartea and Williams (2008)) about appropriate models for gas prices. A standard choice is an Itô diffusion described by a stochastic differential equation (SDE)

$$dG_t = \mu(t, G_t) dt + \sigma(t, G_t) \cdot dW_t, \quad (2)$$

where  $W_t$  is a  $d$ -dimensional Brownian motion and  $\sigma(t, g)$  is a  $d \times d$  non-degenerate volatility matrix. A canonical example (see e.g. Jaillet et al. (2004)) is a one-dimensional exponential Ornstein-Uhlenbeck process, namely

$$\begin{aligned} dG_t &= G_t [\kappa(\theta - \log G_t) dt + \sigma dW_t], \\ \text{or } d(\log G_t) &= \kappa(\theta - \frac{\sigma^2}{2\kappa} - \log G_t) dt + \sigma dW_t, \quad G_0 = g. \end{aligned} \quad (3)$$

This models the mean-reversion (to the average level  $e^\theta$ ) in gas prices documented by Eydeland and Wolyniec (2003), while keeping  $\log G_t$  conditionally Gaussian. Upward jumps in  $(G_t)$  can also be considered and may be used to take into account price spikes. The jury is still out whether such jump-diffusion models are appropriate for natural gas. Other possibilities for  $(G_t)$  could include regime-switching, stochastic mean reversion levels, latent factors, Lévy processes, etc. Our method is independent of the assumed model for  $(G_t)$ , and in general we only make the following technical

### Assumption 1

(A)  $(G_t)$  is a  $d$ -dimensional, strong Markov, non-exploding process in  $\mathbb{R}^d$ .

(B) The information filtration  $\mathbb{F} = (\mathcal{F}_t)$  on the stochastic basis  $(\Omega, \mathcal{F}, \mathbb{P})$  is the natural filtration of  $(G_t)$ .

(C) The reward rate  $\psi_i: [0, T] \times \mathbb{R}^d \times [c_{min}, c_{max}] \rightarrow \mathbb{R}$  is a jointly Lipschitz-continuous function of  $(t, g, c)$  and satisfies

$$\mathbb{E} \left[ \sup_{t \in [0, T]} |\psi_i(t, G_t, C_t)|^2 \middle| G_0 = g, C_0 = c \right] < \infty, \quad \forall g, c, i.$$

For notational clarity we suppress from now on the dependency of  $\psi_i$  and the coefficients of (2) on time  $t$ .

REMARK 2. Above we have stated the model in continuous-time. This is to conform to classical financial stochastic control models; since the final implementation is computer-based and consequently performed in discrete-time, one could also work in discrete-time from the beginning.

### 2.1. Control Problem.

The flexibility available to the manager is specified via the set  $\mathcal{U}$  of possible storage policies  $u$ . For  $t \in [0, T]$ ,  $u_t \in \{-1, 0, 1\}$  denotes the (dynamically chosen) operating regime of the facility. It is convenient to write  $u = (\xi_1, \xi_2, \dots; \tau_1, \tau_2, \dots)$ , where the variables  $\xi_k \in \{-1, 0, 1\}$  denote the sequence of operating regimes taken by  $u$  and  $\tau_k \leq \tau_{k+1} \leq T$  denote the switching times. Thus,  $u_t = \sum_k \xi_k \mathbb{1}_{[\tau_k, \tau_{k+1})}(t)$ , where by convention  $\tau_0 = 0, \xi_0 = i_0$  is the initial facility state.

Given the initial inventory  $C_0 = c$  and the storage strategy  $u$ , the future inventory  $\bar{C}_t(u)$  is completely determined. Namely,  $\bar{C}_t(u)$  satisfies the ordinary differential equation

$$d\bar{C}_s(u) = a_{u_s}(\bar{C}_s(u)) ds, \quad \bar{C}_0(u) = c. \quad (4)$$

In the sequel we will also use the notation  $\bar{C}_t(c, i) \triangleq c + \int_0^t a_i(\bar{C}_s(c, i)) ds$ .

Each change of the facility's regime incurs switching costs. In particular, moving the facility from regime  $i$  to regime  $j$  costs  $K_{i,j} = K(i, j; t, G_t, C_t)$ . This represents both the effort—one must dispatch workers, coordinate with the outgoing pipeline, stop/start the decompressors, etc.—and the time needed to change the operating mode. We assume that the switching costs are discrete:  $K_{i,j} > \epsilon$  for all  $i \neq j$  and some  $\epsilon > 0$ , and  $K_{i,i} = 0$ . For actual salt dome facilities the switching costs are economically negligible; however, in other applications, such as hydro pumped storage, switching costs may be significant. Also, strictly positive switching costs are needed for technical reasons in our continuous-time model in order to guarantee existence of optimal finite switching strategies (i.e. to rule out *chattering*, where the owner would repeatedly change the regimes back-and-forth over a very short amount of time). Since the ultimate computations are in discrete time, switching costs can be set to zero on implementation-level.

A necessary condition for  $u$  to belong to the set  $\mathcal{U}$  of *admissible strategies* is to be  $\mathbb{F}$ -adapted, right-continuous and of  $\mathbb{P}$ -a.s. finite variation on  $[0, T]$ .  $\mathbb{F}$ -adaptiveness is a standard condition implying that the agent only has access to the observed price process and cannot use any other information. Finite variation means that the number of switching decisions must be finite almost surely. Thus,  $\mathbb{P}[\tau_k < T \ \forall k \geq 0] = 0$ . Other restrictions on  $u$  arise from engineering constraints; for example the finite storage constraint requires that  $\bar{C}_t(u) \in [c_{min}, c_{max}]$  for all  $t \leq T$ . Further possibilities are explored in Section 7.2; in the meantime we assume that  $\mathcal{U}(t, c, i)$ , representing the set of all admissible strategies on the time interval  $[t, T]$  starting in regime  $i$  and with initial inventory  $c$ , is a closed subset of  $\mathcal{U}$ .

Subject to those costs and the operational constraints, the facility manager then maximizes the net expected profit. Given initial conditions at time  $t$ ,  $G_t = g, C_t = c$ , and initial operating regime  $i$ , suppose the manager chooses a particular dispatching policy  $u \in \mathcal{U}(t, c, i)$ . If we denote by  $V_u(t, g, c, i)$  the corresponding expected profit until final date  $T$ , then

$$V_u(t, g, c, i) = \mathbb{E} \left[ \int_t^T e^{-r(s-t)} \psi_{u_s}(G_s, \bar{C}_s(u)) ds - \sum_{\tau_k < T} e^{-r\tau_k} K_{u_{\tau_k^-}, u_{\tau_k}} \middle| G_t = g, C_t = c \right]. \quad (5)$$

The first term in (5) counts the total revenues and costs from managing the facility up to the horizon  $T$ , and the second term counts the incurred switching costs. It remains to specify the terminal condition at  $T$ . Typical contracts specify that the facility should be returned with the same inventory  $C_0$  as initially held, and in a certain state, e.g. *store*. To enforce this stipulation, various buy-back provisions are employed. A common condition is

$$V(T, g, c, i; C_0) = -\bar{K}_1 \cdot (C_0 - c)^+ - \bar{K}_2 \cdot (C_0 - c)^- - \bar{K}_3 \mathbb{1}_{i \neq 0}, \quad (6)$$

making the penalty proportional to the difference with stipulated inventory  $C_0$ , with multipliers  $\bar{K}_1$  and  $\bar{K}_2$  used for under-delivery and over-delivery respectively, and adding a second penalty of  $\bar{K}_3$  if the final regime is not *store*. Another common choice is  $V(T, g, c, i; C_0) = -\bar{K}_1 \cdot g \cdot (C_0 - c)^+$ , which penalizes for having less gas than originally and makes the penalty proportional to current price of gas.

Our formal control problem is to compute

$$V(t, g, c, i) \triangleq \sup_{u \in \mathcal{U}(t, c, i)} V_u(t, g, c, i), \quad (7)$$

with  $V(t, g, c, i; u)$  as defined in (5). Besides the value function  $V$  we are also interested in explicitly characterizing an optimal policy  $u^*$  (if one exists) that achieves the supremum in (7). Before

proceeding, let us emphasize the path-dependent nature of (7). Observe that an optimal policy  $u_{t'}^*$  at intermediate time  $t < t' < T$  depends on the current inventory  $C_{t'}$ . However,  $C_{t'} = \bar{C}_{t'}(u^*)$  is itself a function of past strategy  $(u_s^*: t \leq s \leq t')$ . Conversely, current  $C_{t'}$  affects the feasibility of future strategies  $\{u_s, s \geq t'\}$  through the corresponding constraints on  $\mathcal{U}(t', c, i)$ . The standard method of solving control problems is by dynamic programming and would proceed backwards in time, from  $s = T$  towards  $s = t'$ . However, in our case to find the optimal action at time  $t'$  we need to know optimal actions *before*  $t'$ , hence the path-dependency and the resulting challenge. Of course, this was abstractly resolved by making  $C_t$  a state variable in (7). Unfortunately, from a numerical analysis point of view this is only a superficial fix as  $C_t$  is now neither exogenously stochastic, nor directly controlled, creating a degenerate and numerically unstable variable.

### 3. Iterative Optimal Stopping.

Without inventory, (7) belongs to the class of Optimal Switching problems. These have been recently extensively studied, both analytically (Zervos 2003, Pham and Ly Vath 2007, Dayanik and Egami 2005), and numerically (Barrera-Esteve et al. 2006, Porchet et al. 2008). In particular, one can exploit the idea of the authors' earlier paper (Carmona and Ludkovski 2008) to represent (7) as a *sequence* of optimal stopping (American option) problems. These sub-problems precisely capture the timing flexibility of the manager.

Let  $\mathcal{S}_t$  denote the set of all  $\mathbb{F}$ -stopping times between  $t$  and  $T$ . Recursively construct the functions  $V^k(t, g, c, i)$  with  $k = 0, 1, \dots, 0 \leq t \leq T$ ,  $g \in \mathbb{R}^d$ ,  $c \in [c_{min}, c_{max}]$  and  $i \in \{-1, 0, 1\}$  via

$$\begin{aligned} V^0(t, g, c, i) &\triangleq \mathbb{E} \left[ \int_t^T e^{-r(s-t)} \psi_i(s, G_s, \bar{C}_s(c, i)) ds \mid G_t = g \right], \\ V^k(t, g, c, i) &\triangleq \sup_{\tau \in \mathcal{S}_t} \mathbb{E} \left[ \int_t^\tau e^{-r(s-t)} \psi_i(s, G_s, \bar{C}_s(c, i)) ds \right. \\ &\quad \left. + \max_{j \neq i} e^{-r(\tau-t)} \left\{ -K_{i,j} + V^{k-1}(\tau, G_\tau, \bar{C}_\tau(c, i), j) \right\} \mid G_t = g \right]. \end{aligned} \quad (8)$$

Recall that  $\bar{C}_s(c, i)$  denotes the storage level at time  $s$  when policy  $i$  is constantly applied and initial inventory was  $c$ . The results in Carmona and Ludkovski (2008) show that

PROPOSITION 1. *Let  $\mathcal{U}^k(t, c, i) \triangleq \{u \in \mathcal{U}(t, c, i) : u = (\xi_1, \dots, \xi_k; \tau_1, \dots, \tau_k)\}$  be the subset of admissible strategies with at most  $k$  switches. Then*

- i.  $V^k$  is equal to the value function for the storage problem with at most  $k$  switches allowed:  

$$V^k(t, g, c, i) = \sup_{u \in \mathcal{U}^k(t, c, i)} V_u(t, g, c, i).$$

ii. An optimal strategy  $u^* = u^{*,k}$  for  $V^k(0, g, c, i)$  exists, is Markovian and is explicitly defined by  $\tau_0^* = 0, \xi_0^* = i$ , and for  $\ell = 1, \dots, k$  by

$$\begin{cases} \tau_\ell^* \triangleq \inf \left\{ s \geq \tau_{\ell-1}^* : V^\ell(s, G_s, C_s(u^*), i) = \max_{j \neq i} (-K_{i,j} + V^{\ell-1}(s, G_s, C_s(u^*), j)) \right\} \wedge T, \\ \xi_\ell^* \triangleq \arg \max_{j \neq i} \left\{ -K_{i,j} + V^{\ell-1}(\tau_\ell^*, G_{\tau_\ell^*}, C_{\tau_\ell^*}(u^*), i) \right\}. \end{cases} \quad (9)$$

iii.  $\lim_{k \rightarrow \infty} V^k(t, g, c, i) = V(t, g, c, i)$  pointwise, uniformly on compacts.

iv. The limit  $V(t, g, c, i)$  is continuous and is the minimal solution of the Bellman equation

$$V(t, g, c, i) = \sup_{\tau \in \mathcal{S}_t} \mathbb{E} \left[ \int_t^\tau e^{-r(s-t)} \psi_i(G_s, \bar{C}_s(c, i)) ds + e^{-r(\tau-t)} \cdot \max_{j \neq i} \left\{ -K_{i,j} + V(\tau, G_\tau, \bar{C}_{\tau-t}(c, i), j) \right\} \middle| G_t = g \right]. \quad (10)$$

Item (i) says that  $V^k(t, g, c, i)$  is the maximum expected profit to be had on the time period  $[t, T]$  conditional on the initial state  $(g, c, i)$  and at most  $k$  switches remaining. This is useful because according to item (iii), for any  $\epsilon > 0$ , there is a  $K$  large enough such that an optimal control of  $V^K$  as defined in (9), generates an  $\epsilon$ -optimal strategy for  $V$ . The key insight behind the proposition is the Bellman optimality principle which implies that solving the problem with at most  $k+1$  switching decisions allowed is equivalent to finding the first optimal decision time  $\tau$  which maximizes the initial payoff until  $\tau$  plus the value function at  $\tau$  corresponding to optimal switching with  $k$  switches.

REMARK 3. We do not have full results showing the uniqueness or existence of optimal control for the original value function  $V$ , which is a delicate impulse control problem. On a practical level this makes no difference since an  $\epsilon$ -optimal control is always available. Theoretically, it would be interesting to find a good set of working assumptions to ensure optimality existence/uniqueness.

### 3.1. Quasi-variational formulation.

The presented storage model is a special case of stochastic *impulse control* problems. Hence one can apply the generic quasi-variational method developed by Bensoussan and Lions (1984). The verification theorem presented below states that a suitable smooth candidate function  $\varphi$ , which dominates the switching barrier and solves the Kolmogorov pde in the continuation region is indeed the value function of (7). The proof follows from standard techniques, see e.g. Øksendal and Sulem (2005).

PROPOSITION 2. Let  $\mathcal{L}_G$  denote the infinitesimal generator of the Markov process  $(G_t)$ . Suppose there exists  $\varphi(t, g, c, i)$  such that for

$$D \triangleq \bigcup_i \left\{ (t, g, c) : \varphi(t, g, c, i) = \max_{j \neq i} \left\{ -K_{i,j} + \varphi(t, g, c, j) \right\} \right\},$$

$\varphi(\cdot, i)$  belongs to  $\mathcal{C}^{1,2,2}([0, T] \times \mathbb{R}^d \times [c_{\min}, c_{\max}]) \setminus \mathcal{D}) \cap \mathcal{C}^{1,1,1}(\mathcal{D})$  for each  $i \in \{-1, 0, 1\}$  and satisfies the following quasi-variational inequality (QVI):

$$\left\{ \begin{array}{l} \min \left\{ \varphi(t, g, c, i) - \max_{j \neq i} (-K_{i,j} + \varphi(t, g, c, j)), \right. \\ \left. - \frac{\partial}{\partial t} \varphi(t, g, c, i) - \mathcal{L}_G \varphi(t, g, c, i) + a_i(c) \cdot \frac{\partial}{\partial c} \varphi(t, g, c, i) - \psi_i(g, c) + r \varphi(t, g, c, i) \right\} = 0, \\ \varphi(T, g, c, i) = -\bar{K}_1 \cdot (C_0 - c)_+ - \bar{K}_2 \cdot (C_0 - c)_- - \bar{K}_3 \mathbb{1}_{\{i \neq 0\}}. \end{array} \right.$$

Then  $\varphi = V$  is the optimal value function for the storage problem (7).

If the process  $(G_t)$  is an Itô diffusion as in (2), then  $\mathcal{L}_G = \mu(g) \frac{\partial}{\partial g} + \frac{1}{2} \sigma^2(g) \frac{\partial^2}{\partial g^2}$  is a second-order differential operator. The derived parabolic pde system with a free boundary can then be solved using standard tools, see for example (Wilmott et al. 1995, Chapter 7). In the context of gas storage this approach has been explored by Ahn et al. (2002). As the simplest choice, consider the basic finite differencing (FD) algorithm. We set up a uniform space-time grid with steps  $\Delta t$ ,  $\Delta g$  and  $\Delta c$  in the respective variables, and on this grid solve

$$\left\{ \begin{array}{l} 0 = \varphi_t + \mu(g) \varphi_g + \sigma^2(g)/2 \cdot \varphi_{gg} - a_i(c) \cdot \varphi_c + \psi_i(g, c) - r \cdot \varphi, \\ \varphi(t, g, c, i) \geq \max_{j \neq i} (-K_{i,j} + \varphi(t, g, c, j)), \\ \varphi(T, g, c, i) = -\bar{K}_1 \cdot (C_0 - c)_+ - \bar{K}_2 \cdot (C_0 - c)_- - \bar{K}_3 \mathbb{1}_{i \neq 0}, \end{array} \right. \quad (11)$$

by replacing derivatives with explicit finite differences in the first equation and directly enforcing the barrier condition at each time-step. Using standard properties of the infinitesimal generator  $\mathcal{L}_G$ , one obtains the convergence  $\varphi(0, g, c, i) \rightarrow V(0, g, c, i)$  as step sizes vanish  $\Delta t \rightarrow 0$ ,  $\Delta g \rightarrow 0$ ,  $\Delta c \rightarrow 0$ .

The FD method is straightforward to implement but will be slow since even in the easiest case, where  $(G_t)$  is one-dimensional and has smooth dynamics, the pde (11) is *two-dimensional* in space. Furthermore, the degenerate  $C_t$ -dynamics cause numerical instability as the pde is convection-dominated (due to absence of  $\varphi_{cc}$  term). The algorithm is also not robust: for instance, adding jumps to (2) produces a partial *integro-differential* equation which is non-local and requires special numerical tools, see Thompson et al. (2003) for details. Similarly, pde solvers suffer from the curse of dimensionality —making  $(G_t)$  two-dimensional is still beyond today's computational power (in the sense of a business-time system). On the other hand, the error analysis of FD algorithms is well-studied and many improvements are possible, including adaptive solution grids, alternating direction implicit schemes, relaxation methods, etc.

#### 4. Numerical Method.

The benefit of the recursive formulation in (8)-(10) is its suitability for an efficient and scalable numerical implementation. In this section we describe in detail the resulting algorithms.

To begin, we discretize time, setting  $\mathcal{S}^\Delta \triangleq \{m\Delta t, m = 0, 1, \dots, M\}$ ,  $\Delta t = \frac{T}{M}$  as our discrete time grid. Managerial decisions are now allowed only at  $\tau_k \in \mathcal{S}^\Delta$ . This restriction is similar to looking at Bermudan options as approximation to American exercise rights. Denote by

$$\psi_i^\Delta(t, G_t, C_t) \triangleq \int_t^{t+\Delta t} e^{-r(s-t)} \cdot \psi_i(G_s, C_s) ds$$

the total cashflows during one time-step and let  $t_1 = m\Delta t, t_2 = (m+1)\Delta t$  be two generic consecutive time steps. In discrete time, the representation of  $V(t_1, g, c, i)$  in (10) reduces to deciding between immediate switch at  $t_1$  to some other regime  $j$ , which must then be maintained until  $t_2$  (i.e.  $\tau = t_1$  in (10)), versus no switching and therefore maintaining regime  $i$  until  $t_2$  ( $\tau > t_1 \Leftrightarrow \tau \geq t_2$ ). In other words, one chooses the best (in terms of continuation value) regime  $j$  at  $t_1$ , pays the corresponding switching costs, and then waits until  $t_2$ . Thus, using the notation of (4), (8) reduces to

$$V(t_1, G_{t_1}, C_{t_1}, i) = \max_j \left( -K_{i,j} + \mathbb{E} \left[ \psi_j^\Delta(t_1, G_{t_1}, C_{t_1}) + e^{-r\Delta t} \cdot V(t_2, G_{t_2}, \bar{C}_{\Delta t}(C_{t_1}, j), j) \mid \mathcal{F}_{t_1} \right] \right). \quad (12)$$

The dynamic programming method can now be applied to recursively evaluate (12) backwards in time to obtain the discretized Snell envelopes of the optimal stopping problem (10). Hence, for a numeric evaluation of  $V$  it is sufficient to construct an algorithm for evaluating the *conditional expectations* appearing in (12).

Let  $(B_j(g; t_1, c, i))_{j=1}^\infty$  be a given orthonormal basis of  $L^2(\mathcal{F}_{t_1})$  (selection of  $(B_j)$  is discussed below). Recall that  $(G_t)$  is Markov, while  $(C_t)$  is determined by  $u$ . Thus, we may view the conditional expectation in (12) as a map

$$g \mapsto E(g; t_1, c, i) \triangleq \mathbb{E} \left[ \psi_i^\Delta(t_1, G_{t_1}, c) + e^{-r\Delta t} \cdot V(t_2, G_{t_2}, \bar{C}_{\Delta t}(c, i), i) \mid G_{t_1} = g \right]. \quad (13)$$

The latter may be approximated with a projection on the truncated basis  $(B_j)_{j=1}^{N_{b_1}}$ :

$$E(g; t_1, c, i) = \sum_{j=1}^{\infty} \alpha_j B_j(g; t_1, c, i) \simeq \hat{E}(g; t_1, c, i) = \sum_{j=1}^{N_{b_1}} \alpha_j B_j(g; t_1, c, i), \quad (14)$$

where  $\alpha_j$  are the  $\mathbb{R}$ -valued projection coefficients. The right hand side of (14) is a finite-dimensional projection of the continuation values onto the basis functions and can be replaced with an empirical regression based on a Monte Carlo simulation. This then gives a method for implementing (12) on a computer.

We begin by generating  $N$  sample paths  $(g_{m\Delta t}^n)$ ,  $n = 1, \dots, N$  of the discretized  $(G_t)$  process with a fixed initial condition  $G_0 = g = g_0^n$ . As mentioned before, the inventory depends on the policy choice, so it cannot be directly simulated. To overcome this problem, we shall construct a grid in the  $C$ -variable and compute  $V(t, g, c, i)$  only for  $c \in \{c_0 = c_{min}, c_1, \dots, c_{N_C} = c_{max}\}$ .

We will approximate the value function by the empirical average of the pathwise quasi-values (from now on simply values)  $V(0, g, c, i) \simeq \frac{1}{N} \sum_{n=1}^N v(0, g_0^n, c, i)$ . The values  $v(t, g_t^n, c, i)$  are computed recursively in a backward fashion, starting with the terminal condition of (6):  $v(T, g_T^n, c, i) = -\bar{K}_1 \cdot (C_0 - c)_+ - \bar{K}_2 \cdot (C_0 - c)_- - \bar{K}_3 \mathbb{1}_{i \neq 0}$ . Consider again two consecutive time steps  $t_1, t_2$  and suppose inductively that we know  $v(t_2, g_{t_2}^n, c, i)$  along the paths  $(g_{t_2}^n)_{n=1}^N$  and for  $c \in \{c_\ell, \ell = 1, \dots, N_C\}$ . Our goal is to compute  $v(t_1, g_{t_1}^n, c, i)$ . To obtain the prediction  $\hat{E}(g_{m\Delta t}^n; t_1, c, i)$  of the continuation value, one first computes  $v(t_2, g_{t_2}^n, \bar{C}_{\Delta t}(c, i), i)$ . Note that in general  $\bar{C}_{\Delta t}(c, i)$  does not belong to the grid  $\{c_\ell\}$ , and interpolation is needed. Then one regresses  $v(t_2, g_{t_2}^n, \bar{C}_{\Delta t}(c, i), i)$  against the basis functions  $(B_j(g_{t_1}^n; t_1, c, i))_{j=1}^{N_{b_1}}$  to find the corresponding  $\alpha_j \equiv \alpha_j(t_1, c, i)$  and applies (14). By analogue of (12), the estimate for  $v(t_1, g_{t_1}^n, c, i)$  is then

$$v(t_1, g_{t_1}^n, c, i) = \hat{E}(g_{t_1}^n; t_1, c, i) \vee \max_{j \neq i} \left( -K_{i,j} + \hat{E}(g_{t_1}^n; t_1, c, j) \right). \quad (15)$$

Observe that (15) performs *pathwise* computations, while using across-the-paths projection  $\hat{E}$ . The scheme (15) first appeared in Tsitsiklis and van Roy (2001) in the context of American option pricing. In our setting we call it a mixed interpolation-Tsitsiklis-van Roy scheme (MITvR).

It is also useful to think in terms of the optimal storage strategy. Let  $\hat{j}^n(t_1; i) \in \{-1, 0, 1\}$  represent the optimal decision on the  $n$ -th path at time  $t = t_1$  and current regime  $i$ . The analogue of (12) implies that (recall  $K_{i,i} \equiv 0$ )

$$\hat{j}^n(t_1; i) = \arg \max_j \left( -K_{i,j} + \hat{E}(g_{t_1}^n; t_1, c, j) \right). \quad (16)$$

Thus, the set of paths on which it is optimal to switch at time  $t = m\Delta t$  is given by  $\{n: \hat{j}^n(m\Delta t; i) \neq i\}$ . Equation (16) can be used to construct the switching boundaries which partition  $[0, \infty] \times [c_{min}, c_{max}]$  into regions where injection, withdrawal, etc., is optimal at date  $t$ .

The efficiency of (15) is enhanced by using *the same set* of paths to compute all the conditional expectations. Nevertheless, because of the capacity variable  $C$  the above approach is still time-intensive. Indeed, at every time step  $m\Delta t$  and regime  $i$ , we must run a separate regression for each inventory grid point  $c_\ell$ . Hence, in terms of computational complexity, the above method is equivalent to solving  $N^c$  optimal switching problems.

The choice of appropriate basis functions  $(B_j(\cdot; t, c, i))$  in (14) is user-defined. A detailed analysis of different orthogonal families is available in Stentoft (2004). Empirically, basis choice has only a mild effect on numerical precision, but strongly affects the *variance* of the algorithm. Thus, customization is desirable and it helps to use basis functions that resemble the expected shape of the value function. In practice,  $N_{b_1}$  as small as five or six normally suffices, and having more bases

can often lead to worse numerical results due to overfitting. Let us mention that the requirement of an orthonormal basis is purely theoretical and any set of linearly independent functions will suffice. Some of our favorite choices are exponential functions  $e^{\alpha g}$  and the polynomials  $g^m$ ; this choice is essentially heuristic. Also, we typically select the bases independent of parameters  $(t, c, i)$ , though the latter offer a wide scope for additional fine-tuning.

#### 4.1. Quasi-Simulation of Inventory Levels.

To maintain numerical efficiency it is desirable to avoid the fixed discretization in the  $C$ -variable that resembles the slow lattice schemes. Accordingly, we propose the following alternative that uses *pathwise* and regime-dependent inventory levels  $(c_{m\Delta t}^n(i))$ . The idea is to perform a bivariate regression in (14) of tomorrow's value against the (price, inventory) pair. The paths  $(c_{m\Delta t}^n(i))_{m=1}^M$  are generated backwards during the dynamic programming procedure by combining randomization and guesses of today's optimal strategy. Besides added efficiency, we are also guided by considerations of *accuracy*. Quasi-simulation of inventory allows us to use the Longstaff and Schwartz (2001) scheme of computing pathwise value functions of optimal stopping problems. From simpler problems of American option pricing and plain optimal switching we know that the LSM scheme typically has less bias (though more variance) than the TvR scheme (15) (Ludkovski 2005).

We inductively assume again that we are given the  $3N$  values  $v(t_2, g_{t_2}^n, c_{t_2}^n(i), i)$ ,  $i \in \{-1, 0, 1\}$ ,  $n = 1, \dots, N$ , as well as *bivariate* basis functions  $(\bar{B}_j(g, c; t_1, i))_{j=1}^{N_{b_2}}$ . For a given path  $n$ , regime  $i$ , and a given inventory  $c_{t_1}^n(i)$  (see below about obtaining  $c_{t_1}^n(i)$ ) we make the optimal switching decision as follows (compare with (15)):

1. For each  $k \in \{-1, 0, 1\}$ , regress  $\{e^{-r\Delta t} \cdot v(t_2, g_{t_2}^n, c_{t_2}^n(k), k)\}_{n=1}^N$  against the basis functions  $(\bar{B}_j(g_{t_1}^n, c_{t_2}^n(k); t_1, k))_{j=1, \dots, N_{b_2}}$ . This gives a prediction

$$\tilde{E} : (g, c, k) \mapsto \sum_{j=1}^{N_{b_2}} \bar{\alpha}_j \bar{B}_j(g, c; t_1, k) \simeq \mathbb{E} \left[ \psi_k^\Delta(t_1, g, \bar{C}_{-\Delta t}(c, k)) + e^{-r\Delta t} \cdot v(t_2, G_{t_2}, c, k) \mid G_{m\Delta t} = g \right] \quad (17)$$

of the value tomorrow given today's prices and *tomorrow's* inventory.

2. Compute  $\bar{C}_{\Delta t}(c_{t_1}^n(i), j)$ , the inventory tomorrow given today's inventory  $c_{t_1}^n(i)$  and the decision to switch to  $j$ .

3. The optimal decision is the regime  $\hat{j}^n(t_1, i)$  maximizing the approximate continuation value, cf. (12),

$$\hat{j}^n(t_1, i) = \arg \max_j \left\{ \tilde{E}(g_{t_1}^n, \bar{C}_{\Delta t}(c_{t_1}^n(i), j), j) - K_{i,j} \right\}. \quad (18)$$

4. If  $\bar{C}_{\Delta t}(c_{t_1}^n(i), \hat{j}^n) = c_{t_2}^n(\hat{j}^n)$ , then the Longstaff-Schwartz update is used:

$$v(t, g_{t_1}^n, c_{t_1}^n(i), i) = \int_{t_1}^{t_2} e^{-r(t-t_1)} \psi_{\hat{j}^n}(G_t, \bar{C}_{t-t_1}(c_{t_1}^n(i), \hat{j}^n)) dt + e^{-r\Delta t} \cdot v(t_2, g_{t_2}^n, c_{t_2}^n(\hat{j}^n), \hat{j}^n) - K_{i, \hat{j}^n}. \quad (\text{LSM})$$

Else, one updates via

$$v(t_1, g_{t_1}^n, c_{t_1}^n(i), i) = \tilde{E}(g_{t_1}^n, \bar{C}_{\Delta t}(c_{t_1}^n(i), \hat{j}), \hat{j}) - K_{i, \hat{j}}. \quad (\text{TvR})$$

The first case (LSM) stands for Least Squares Monte Carlo. In that version the across-the-paths regression is used primarily to make the optimal switching decision, but is not necessarily fed into the pathwise values. This helps to eliminate potential biases from the regressions by preventing error accumulation across time-steps. To preserve this beneficial look-ahead property of the Longstaff and Schwartz (2001) algorithm, we attempt to *speculatively* pick  $c_{m\Delta t}^n(i)$  such that the first case (LSM) occurs as much as possible. In other words, as we move back in time, we try to select inventory levels that form an optimal (price, inventory)-path on the remaining time interval. When this is not possible (due to capacity or other constraints, or if our guess of  $\hat{j}$  is incorrect), we fall back onto the basic (TvR) scheme. Accurate guessing of  $\hat{j}$  means that we correctly select the optimal strategy (up to the errors resulting from the projection). In such a case, we have  $v(t_1, g_{t_1}^n, c_{t_1}^n(i), i) = \int_{t_1}^T e^{-r(s-t)} \psi_{u_s^*}(G_t, \bar{C}_s(u_s^*)) ds - \sum_{\tau_k^* < T} e^{-r(\tau_k - t_1)} K_{u_{\tau_k^*}^*, u_{\tau_k^*}^*}$  exactly along the price path. Observe also that the method can be used iteratively over several simulation runs, improving the guesses of  $\hat{j}$  over time.

The terminal inventory levels  $c_T^n$  are randomized and obtained by independent and uniform samples from  $[c_{min}, c_{max}]$ . At each step  $t = m\Delta t$ , some randomization in  $(c_{t_1}^n(i))$  is also desirable in order to avoid clustering and allow for good fit during the regression step. Of course, randomization reduces the number of paths satisfying (LSM), and balancing the two objectives is a detail that we leave to implementation. We christen this scheme Bivariate Least Squares Monte Carlo (BLSM).

## 4.2. Algorithm Summary

1. Select a set of univariate basis functions  $(B_j)$ , bivariate basis functions  $(\bar{B}_j)$  and algorithm parameters  $\Delta t, M, N, N_{b_1}, N_{b_2}$ .

2. Generate  $N$  paths of the price process:  $\{g_{m\Delta t}^n, m = 0, 1, \dots, M, n = 1, 2, \dots, N\}$  with fixed initial condition  $g_0^n = g_0$ . Generate a random terminal inventory level  $c_T^n(i)$  for each path and each regime  $i$ .

3. Initialize the pathwise values  $v(T, g_T^n, c_T^n(i), i)$  from (6).

4. Moving backward in time with  $t = m\Delta t$ ,  $m = M, \dots, 0$  repeat the Loop, where the computations are based on (10):

*i) Guess Current Inventory:* generate  $(c_{m\Delta t}^n(i))$  by guessing the optimal decision  $\hat{j}^n(m\Delta t, i)$  and solving  $c_{m\Delta t}^n(i) = \bar{C}_{-\Delta t}(c_{(m+1)\Delta t}^n(\hat{j}^n(m\Delta t, i), \hat{j}^n(m\Delta t, i)))$ .

*ii) Regression Step:* do the bivariate regression of (17).

*iii) Optimal Decision Step:* find the optimal decision using (18).

*iv) Update Step:* compute  $v(m\Delta t, g_{m\Delta t}^n, c_{m\Delta t}^n(i), i)$  via (LSM) and (TvR).

*v) Switching Sets:* the points

$$\mathcal{C}_{m\Delta t}(i, j) \triangleq \{(g_{m\Delta t}^n, c_{m\Delta t}^n) : n \text{ is such that } \hat{j}^n(m\Delta t, i) = i\} \quad (19)$$

define the empirical region in the  $(G, C)$ -space where switching from regime  $i$  to regime  $j$  is optimal.

This defines the optimal strategy at  $t = m\Delta t$ .

5. end Loop

6. Interpolate  $V(0, g_0, c, i)$  from the  $N$  values  $v(0, g_0^n, c_0^n(i), i)$  for the desired inventory level  $c$  (using splines, kernel regression, etc.).

REMARK 4. As mentioned before, in the discrete-time version we allow switching costs to be zero,  $K_{i,j} \equiv 0$ . In that case  $V(t, g, c, i)$  does not depend on the current regime  $i$  and so one can save on the corresponding computations.

### 4.3. Algorithm Complexity.

The BLSM algorithm requires  $\mathcal{O}(N \cdot M \cdot ((N_{b_1})^3 + (N_{b_2})^3))$  operations. The most computationally intensive operation is the regression step where we face matrices of size  $N \times N_{b_1}$  and  $N \times N_{b_2}$ , and which make the algorithm linear in the larger dimension  $N$  and cubic in the smaller dimensions  $N_{b_1}, N_{b_2}$ . In contrast, the algorithm complexity of the MITvR scheme is  $\mathcal{O}(N \cdot M \cdot N_c \cdot (N_{b_1})^3)$  where  $N_c$  is the grid size in the  $C$ -variable. However, these expressions hide the relationship between  $N_{b_1}, N_{b_2}$  and  $N$ , because more basis functions require more paths for accurate evaluation of the regression step. In fact, according to Glasserman and Yu (2004),  $N$  must be asymptotically exponential in the number of basis functions. On the other hand, to perform the bivariate regression (17), it is likely that a large number of basis functions  $N_{b_2}$  is needed, about 12 – 15 in our experience. Hence in the BLSM algorithm  $N$  must be taken larger than in the the MITvR case. Precise comparison is hard because the BLSM scheme inherently generates more variance and we have no hard benchmark to go by. For our examples we find that  $N = 16,000$  for the MITvR scheme is reasonable, while  $N = 40,000$  is needed for BLSM. Practically speaking this implies that BLSM is about twice as fast as MITvR, see Section 5. The memory requirements of both schemes are  $\mathcal{O}(N \cdot M)$  corresponding to the need to store the entire sample paths  $(g_{m\Delta t}^n)_{n=1}^N$  in memory.

#### 4.4. Convergence.

The presented algorithms have several layers of approximations. Three major types of errors can be identified: error due to time discretization and the corresponding restriction of strategies to  $\mathcal{U}^\Delta$ , projection error and Monte Carlo sampling error. Detailed error analysis has been performed in Carmona and Ludkovski (2008) for the case of the TvR scheme with no inventory. Taking  $C_t$  to be a ‘dummy’ variable determined by the dynamics of  $G_t$  and policy  $u$ , the results carry over without change. Analysis of the BLSM scheme is too involved, however see partial results in this direction in Egloff (2005). This lack of provable convergence results is typical for Monte Carlo optimal stopping methods, largely due to the nonlinearity introduced by the stopping boundaries. Nevertheless, extensive empirical experiments (Stentoft 2004) have strongly supported the general TvR/LSM methodology.

The error from discretizing ( $G_t$ ) and simultaneously restricting the switching times to occur only at the discrete time grid points is known to be  $\mathcal{O}(\sqrt{\Delta t})$ . The error from approximating the conditional expectations with a projection (14) is on the order of  $\mathcal{O}(\Delta t^{-k} \cdot (\|E - \hat{E}\|))$  when computing  $V^k$  (Carmona and Ludkovski 2008). This suggests that the projection errors multiply in the number of decisions taken  $k$ . However, empirically the dependence on  $\Delta t$  is much better, especially under the BLSM scheme, so this upper bound is probably not tight. In any case for practical examples, the typical number of switches is in single digits. Finally, the third source of error is due to approximating the projections with an empirical regression using  $N$  realizations of the paths ( $g_{m\Delta t}^n$ ,  $n = 1, \dots, N$ ). This error is difficult to analyze due to interactions between the path-by-path maximum taken in (16) and the across-the-paths regression. No convergence behavior is known; however numerical experiments suggest that it is close to  $\mathcal{O}((\Delta t \cdot N)^{-1/2})$ , which is the expected rate for Monte Carlo methods. Table 1 illustrates this conjecture on Example 2 below. We run the BLSM algorithm using 8000 – 40000 Monte Carlo paths and tabulate the resulting standard errors. At least in this case we see in fact a faster than  $\mathcal{O}(N^{-1/2})$  convergence. **Some residual variance in  $V$  remains even for large  $N$  due to the sensitivity of the algorithm to outliers in generated paths of  $g_{m\Delta t}^n$  for small  $m$ . Since all paths begin from the same initial value  $g_0^n$ , the regression step is sensitive in the first few time steps.**

## 5. Numerical Results.

In this section we present several examples to show the structure of the storage problem and the scalability of our algorithm.

**Table 1** Convergence of Monte Carlo error for Example 1 under the BLSM scheme. The initial gas price is  $G_0 = 3$  \$/MMBtu, initial inventory is  $C_0 = 4$  Bcf, initial regime is  $i_0 = store$  and the horizon is  $T = 1$  year with  $\Delta t = 0.005$ . The table shows the mean and standard deviation of  $V(0, G_0, C_0, i_0)$  (in MM\$/MMBtu) obtained by running the algorithm 50 times.

Paths $N$	Mean	Std. Dev
8000	9.63	0.4179
16000	9.41	0.1362
24000	9.37	0.0961
32000	9.38	0.0663
40000	9.35	0.0647

**Table 2** Comparison of numerical results for Example 1. The initial gas price is  $G_0 = 3$  \$/MMBtu, initial inventory is  $C_0 = 4$  Bcf, initial regime is  $i_0 = store$  and the horizon is  $T = 1$  year. The table shows  $V(0, G_0, C_0, i_0)$  (in MM\$/MMBtu) computed using four different methods. Standard deviations for the Monte Carlo methods were obtained by running the respective algorithms 50 times.

Method	$V(0, G_0, C_0, i_0)$	Std. Dev	Time (min)
Coarse FD	9.32	–	24
Fine FD	9.44	–	65
MITvR	9.86	0.021	47
BLSM	9.35	0.067	32

EXAMPLE 1. As a first illustration of our approach, consider a facility with a total capacity of  $c_{max} = 8$  Bcf rented out for one year,  $T = 1$ . The price process is taken from the data of de Jong and Walet (2003),

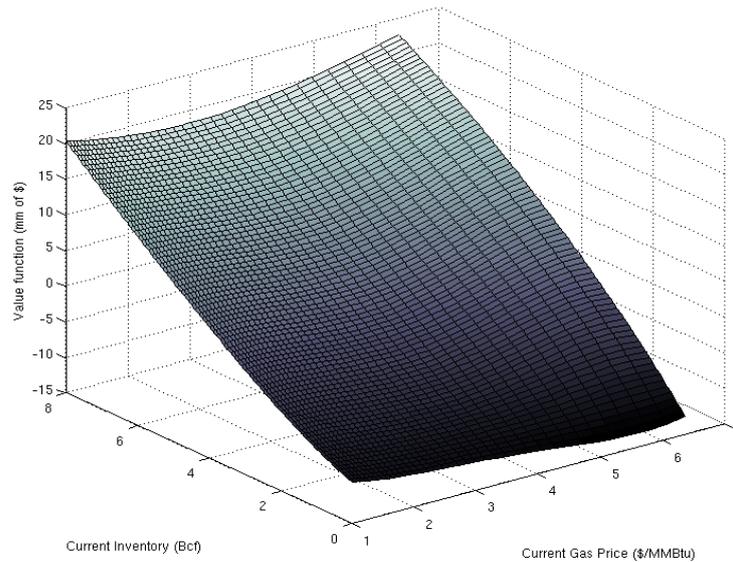
$$d \log G_t = 17.1 \cdot (\log 3 - \log G_t) dt + 1.33 dW_t. \quad (20)$$

Observe the very fast mean-reversion  $\kappa = 17.1$  of the prices, with a half-life of 15 days. The initial inventory is 4Bcf and the terminal condition is  $V(T, g, c, i) = -2 \cdot g \cdot \max(4 - c, 0)$ . Thus, the manager is penalized at double the market price for final inventory being less than 4 Bcf and receives no compensation for any excess. The other parameters (in yearly units) in (1) are

$$\left\{ \begin{array}{l} a_{in}(c) = 0.06 \cdot 365, \quad K_i(c) \equiv 0.1c, \quad K_{i,j} \equiv 0.25 \text{ for } i \neq j, \\ a_{out}(c) = 0.25 \cdot 365, \quad r = 0.06, \quad b_i(c) \equiv a_i(c). \end{array} \right\}$$

Thus, it takes about  $8/0.06 = 133$  days to fill the facility and  $8/0.25 = 32$  days to empty it. In this simple example we have taken the injection/withdrawal rates to be independent of inventory levels and assumed lossless pumping.

We solve this storage problem using three different solvers: an explicit finite-difference pde solver discretizing (11), the MITvR scheme of (15) and the BLSM scheme of (LSM). The results are



**Figure 1** Value function surface for Example 1 showing  $V(0.5, g, c, store; T = 1)$  as a function of current gas price  $G_t = g$  and current inventory  $C_t = c$ .

summarized in Table 2. As an extra check we used two different grid sizes for the pde solver: a coarse  $250 \times 250$   $(G, C)$ -grid with 10000 time steps and a finer  $500 \times 500$   $(G, C)$ -grid with 20000 time steps. The MITvR scheme used 200 time steps, 10000 paths with six basis functions and 80 grid points in the  $C$ -variable. The quasi-simulation BLSM scheme used 200 time steps and 40,000 paths with fifteen basis functions.

Taking the fine pde solver as the benchmark value, we see that the simulation methods are within 5% of the optimal value and seem to have an upper bias. **In particular, the BLSM algorithm with 40,000 paths produces a storage value that is within 1% of the fine FD value, giving an acceptable level of accuracy.** The computational challenges involved are indicated by the long running times of the algorithms.<sup>4</sup> In this light, the 45% time savings obtained by the joint  $(G, C)$ -regression are significant from a practical point of view.

Figure 1 shows the value function  $V(t, g, c, i)$  as a function of current price and inventory for an intermediate time  $t = 0.5$  and  $i = store$  regime. Not surprisingly, higher inventory increases the value function since one has the opportunity to simply sell the extra gas on the market. In the  $G_t$ -variable we observe a parabolic shape with a minimum around the long-term mean 3\$/MMBtu. Thus, deviations of  $G_t$  from its mean imply higher future profits, confirming our intuition about storage acting as a financial straddle.

<sup>4</sup>The simulation methods were run in Matlab on a 1.6GHz desktop. The pde solver was written in C++ and run on the same machine.

We next investigate the sensitivities of the value function. Table 3 shows the effect of storage flexibility on  $V$ . Higher transmission rates increase the extrinsic value of storage, since the manager can move more gas in and out of the facility under “favorable” circumstances. In the example considered, the smaller injection rate acts as a bottleneck on the manager’s flexibility, so the derived extrinsic value is more sensitive to  $a_{inj}$  than to  $a_{out}$ . Table 4 studies the effect of other parameters on the extrinsic value. We find that switching costs  $K_{i,j}$  have a major impact on  $V$ . High  $K_{i,j}$  makes the manager risk-averse and unwilling to change the pumping regime until a very good opportunity comes along (since each switch has a high upfront fixed cost). We also find that because of the limited transmission rates and the mean-reverting nature of the prices, there are dis-economies of scale with respect to facility size. Thus, cutting the facility size to 6Bcf reduces value by nearly 16%, but an increase from 10Bcf to 12Bcf produces a benefit of just 3%. **The last panel of Table 4 studies the impact of price mean-reversion on value of storage. Intuition suggests that stronger mean-reversion would reduce price fluctuations, and therefore also the opportunities to capitalize on temporal price spreads. This is illustrated in Table 4 which shows that increasing the mean-reversion strength  $\kappa$  in (20) decreases storage value.**

**Table 3** Effect of Storage Flexibility on the Value Function in Example 1. Extrinsic value corresponds to  $V(0, 3, 4, 0)$ . Results obtained using the BLSM algorithm with  $N = 40000$  paths and  $\Delta t = 0.005$ .

$a_{in}$ Bcf/day	$a_{out}$ Bcf/day	Extrinsic Value
0.06	0.25	9.35
0.03	0.125	4.75
0.12	0.50	14.50
0.18	0.75	17.28
0.12	0.25	12.33

EXAMPLE 2. Our second example is based on the situation presented in Thompson et al. (2003). A mean-reverting model is taken for gas prices, with a seasonally-adjusted mean-reverting level. The gas prices satisfy

$$dG_t = 4 \cdot (6 + \sin(4\pi t) - G_t) dt + 0.5G_t dW_t.$$

Thus, the mean-reversion level has a seasonal component representing the summer/winter price increases in the North American markets. This seasonality implies an approximate trough-to-peak calendar spread of \$1 for each half year.

**Table 4** Effect of Engineering Characteristics on the Value Function in Example 1. Results obtained using the BLSM algorithm with 40,000 paths.

Effect of Facility Size	
Capacity (Bcf)	$V(0, 3, 4, 0)$
6	7.78
8	9.35
10	10.26
12	10.58
Effect of Switching Costs	
$K_{i,j}, i \neq j$	$V(0, 3, 4, 0)$
0.01	13.25
0.1	11.40
0.25	9.35
0.5	6.73
Effect of Storage Costs	
$K_i$	$V(0, 3, 4, 0)$
0	9.77
$0.05 \cdot c$	9.56
$0.1 \cdot c$	9.35
Effect of Mean Reversion Strength	
$\kappa$	$V(0, 3, 4, 0)$
7.1	10.47
12.1	10.32
17.1	9.35
22.1	8.73

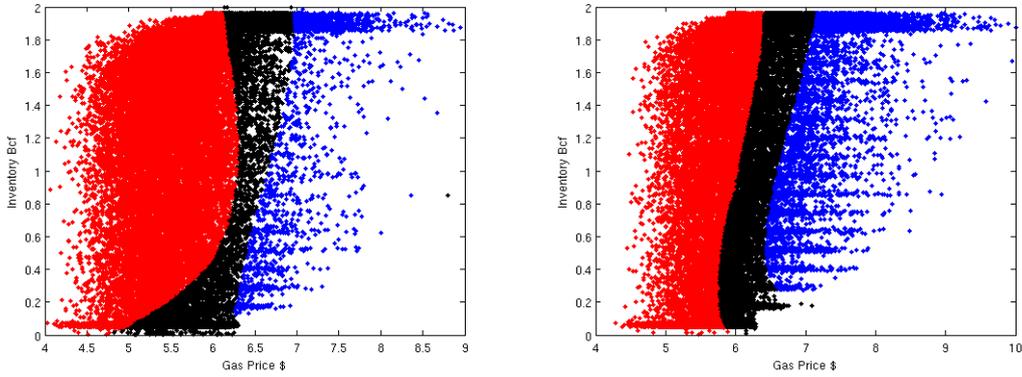
Secondly, the injection and withdrawal rates are ratcheted in terms of current inventory. Thus, injection rate decreases as the amount of gas in the facility grows, and conversely withdrawal rate decreases as less gas is on inventory. More precisely, the facility capacity is  $c_{max} = 2$  Bcf and the yearly rates are  $a_{out}(c) = 0.177 \cdot 365\sqrt{c}$ ,  $a_{in}(c) = 0.0632 \cdot 365\sqrt{\frac{1}{c+0.5} - \frac{1}{2.5}}$ . These rate functions are related to the ideal gas law. Namely, gas transmission rate is proportional to pressure in the reservoir, which in turn is inversely quadratically related to gas volume. Thus, maximum injectivity at  $c_{min} = 0$  is 0.08Bcf/day, and maximum withdrawal at  $c_{max} = 2$  is 0.25Bcf/day.

The monetary reward functions are given by

$$\begin{cases} \text{Inject:} & \psi_{-1}(g, c) = -(a_{in}(c) + 0.6205)g, & dC_t = (a_{in}(C_t)) dt, \\ \text{Store:} & \psi_0 = 0, & dC_t = 0 dt, \\ \text{Withdraw:} & \psi_1(g, c) = a_{out}(c)g, & dC_t = -a_{out}(C_t) dt. \end{cases} \quad (21)$$

Note that there is gas loss during injection represented by the term  $0.6205g$ . We take a horizon of one year  $T = 1$  with no terminal penalty,  $V(T, g, c, i) = 0$ . There are no switching costs and  $r = 0.1$ . Figure 2 presents the optimal control for Example 2 at different times to maturity. The three shades indicate the regions of injection (on the left), storage (dark region in the middle) and

withdrawal (on the right) respectively. We observe that the optimal action regions  $\mathcal{C}_{m\Delta t}(i, j)$  of (19) are strongly dependent on time to maturity  $t$ . Namely, as the terminal horizon approaches, the no-action regions widen, as the fixed switching costs begin to dominate the potential benefit from changing regimes. Also, the terminal condition starts to influence the manager's optimal decision. Figure 2 also shows that  $\mathcal{C}_{m\Delta t}(i, j)$  depends on  $i$ , since the positive switching costs induce inertia and make the manager reluctant to change policies. Finally, the ratcheting transmission rates cause the curvatures in  $\mathcal{C}_{m\Delta t}(i, j)$  as a function of current inventory  $c$ .



Switching regions  $\mathcal{C}_{0.25}(0, \cdot)$  from regime *store* at 3 months to maturity.      Switching regions  $\mathcal{C}_{0.5}(1, \cdot)$  from regime *withdraw* at 6 months to maturity.

**Figure 2** Optimal Controls for Example 2 using the BLSM algorithm with 32,000 paths. Each point corresponds to a simulated  $(g_t^n, c_t^n(i))$  pair, and the color indicates the optimal  $\hat{j}$  of (18) (red for *inject*, black for *store*, blue for *withdraw*).

EXAMPLE 3. Finally, in our third example we illustrate the flexibility of the simulation approach with regards to more complex price processes. It is well known that a one-factor diffusive model does not provide a good fit to gas prices. Accordingly, let us consider a two-factor model with jumps; namely a log-mean-reverting diffusive factor and a second mean-reverting pure jump factor. The second factor captures spikes in natural gas prices without making the mean-reversion rate unnecessarily high (Kluge (2004)):

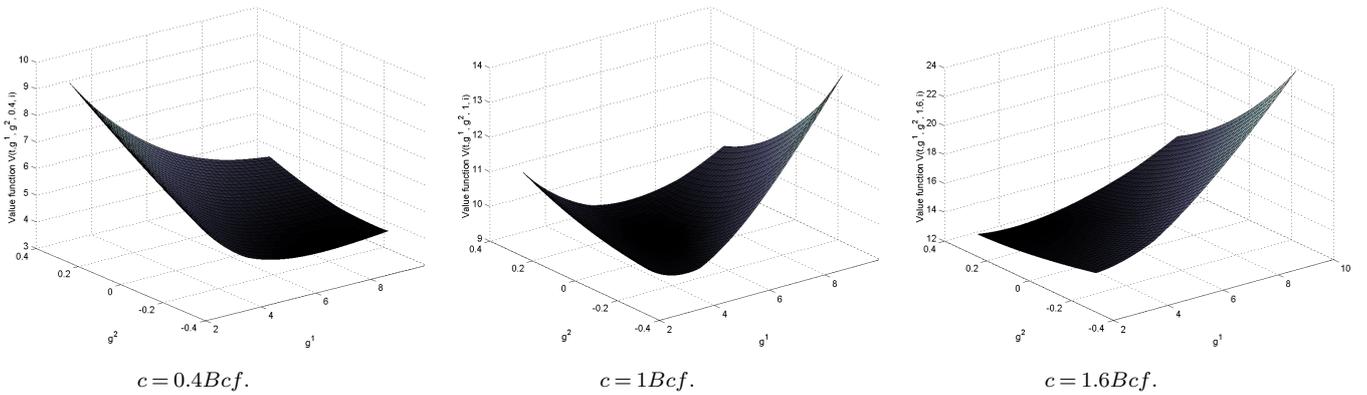
$$\begin{cases} dG_t^1 = 4(\log 6 - \log G_t^1)G_t^1 dt + 0.5G_t^1 dW_t, \\ dG_t^2 = 26(0 - G_t^2) dt + \xi_t dN_t - \lambda\mu_J dt, \end{cases} \quad (22)$$

where  $(N_t)$  is an independent Poisson process with intensity  $\lambda$ , and the jump size  $\xi_t \sim \mathcal{N}(\mu_J, \sigma_J)$  has normal distribution. The total gas price is the product  $G_t \triangleq G_t^1 \cdot \exp(G_t^2)$ , and  $G^2$  can be interpreted as the multiplicative jump factor that causes price spikes on the scale of  $\mu_J\%$ . The possibility of price spikes makes storage much more valuable since it increases the volatility of

inter-temporal spreads. We pick  $\lambda = 12, \mu_J = 0.02, \sigma_J = 0.1$  for the jump component, as well as  $T = 2, r = 0.06, K_{i,j} = K_i \equiv 0, V(T, \cdot) = 0$ .

Implementing Example 3 requires only minor modifications to the implementation of Example 2, which essentially reduce to simulation of the bivariate price process  $(G_t^1, G_t^2)$  and selection of bivariate basis functions  $\bar{B}_j(g^1, g^2)$ . These changes only take a few minutes to make, and the resulting algorithm will take only a little longer to run (depending on how many basis functions are added to deal with  $(G_t^2)$ ). In contrast, with a pde approach, the new state dimension would require an extensive rewrite of the code and would slow the performance by an order of magnitude.

Since the value function  $V(t, g_1, g_2, c, i)$  now has three space variables, in Figure 3 we visualize the dependence of  $V$  just on the two price factors  $(g_1, g_2)$  for different inventory levels  $c$ . Since each factor is mean-reverting and the total price is a product of the two, the value function will exhibit a parabolic straddle shape in each factor. Thus, in Figure 3, when  $g_2 < 0$  one can expect prices to rise back to their “normal” level, and so this presents an opportunity for injection, at least when inventory is low. Conversely, when  $g_2 > 0$  and inventory is high, we are in an upward spike with prices expected to fall and an attractive withdrawal opportunity. The dependence of  $V$  on  $g_1$  is similar to that of Figure 1. Overall, as a function of the price and the spike factor, the value function exhibits an asymmetric “bowl” shape, which in turn is dependent on the current level of inventory.



**Figure 3** Value function  $V(1, (g^1, g^2), c, -1)$  for Example 3 for different inventory levels  $c$ .

## 6. Hydroelectric Pumped Storage.

Another important practical application of our model is hydroelectric pumped storage. Pumped storage consists of a large reservoir of water held by a hydroelectric dam at a higher elevation.

When desired, the dam can be opened which activates the turbines and moves the water to another, lower reservoir. The generated electricity is sold to the power grid. As the water flows, the upper reservoir is depleted. Conversely, in times of low electricity demand (weekends, etc.), the water can be pumped back into the upper reservoir using special electricity-operated pumps (with the required energy purchased from the grid). The efficiency of the system is about 80%, and the capacity of such pumped storage facilities is typically on the order of several hundred megawatt-hours (MWh). Currently pumped storage is the dominant type of electricity storage with more than a hundred facilities around the world (ASCE 1996). **The major use of pumped storage is to capture the daily off-peak/on-peak spread in electricity prices. However, by adjusting pump rates, the reservoir manager can also take advantage of seasonal price/water level fluctuations like in the gas storage problem, see Blöchlinger et al. (2004).**

Beyond direct losses from upstream pumping, stored water is subject to evaporation. At the same time, precipitation and/or upper river run-off provide reservoir replenishment. Realistic modeling is complicated by the need to compute the potential energies of the reservoirs which depend on the relative levels of the water and in turn modify generation rates  $a_i(C_t)$ . We abstract from these concerns and treat the problem in our framework of commodity storage (5), with an addition of another variable modeling weather. Let  $L_t$  represent the Markovian weather state at time  $t$  (e.g.  $L_t$  can be a humidity index or river flow rate).  $(L_t)$  controls reservoir gains/losses, so that inventory depletes at rate  $d(L_t, C_t)$  irrespective of the storage regime. The inventory  $C_t$  represents water level in the upper reservoir<sup>5</sup>. The pumping inefficiency is represented by a multiplier  $\bar{K} > 1$ ,  $b_{in} = \bar{K}a_{in}$ ,  $b_{out} = a_{out}$  that affects the cost of pumping. The overall model is thus:

$$\begin{cases} \text{Pump:} & \psi_{-1}(g, c) = -\bar{K} \cdot g \cdot a_{in} - K_{-1}(c), & dC_t = [a_{in} - d(L_t, C_t)] dt, \\ \text{Store:} & \psi_0(g, c) = -K_0(c), & dC_t = [a_0 - d(L_t, C_t)] dt, \\ \text{Generate:} & \psi_1(g, c) = +g \cdot a_{out} - K_1(c), & dC_t = [-a_{out} - d(L_t, C_t)] dt. \end{cases} \quad (23)$$

Once a suitable model is chosen for  $(L_t)$  (see e.g. Cao et al. (2004)), implementation would be similar to Example 2 above, and would require minimal changes to the simulation algorithm. Note that one could mix-and-match different model types for different variables, for instance a jump-diffusion model for gas prices, and a seasonal  $AR(1)$  model for river run-off. Such flexibility would be hard to achieve outside of simulation paradigm (compare to the stochastic programming approach of Nowak and Römisich (2000), Doege et al. (2006)).

<sup>5</sup> A full model should also specify the lower reservoir inventory since the latter also depletes over time.

## 7. Extensions.

In this section we discuss various extensions and modifications that can be made to our model. First, let us remark that many other resource management problems can be recast in our framework. Such problems all feature fluctuating commodity prices, finite inventory constraints and a small number of operating regimes describing the facility state. Below we elaborate on some of the possibilities.

### 7.1. Other Applications.

**7.1.1. Mine management** A producer extracts metal from a mine with initial capacity  $C_0$ . As the resource is mined, the inventory  $C_t$  declines. In the meantime, the producer can control the mine operating regime to time the extraction with high commodity prices represented by  $(G_t)$ . In this situation, the remaining resource amount  $C_t$  is non-decreasing, since only extraction is possible. Exhaustion implies that no profit is available when  $C_t = 0$ :  $V(t, g, 0, i) = 0$ . Armed with our methodology we can e.g. redo in a more efficient manner (see Ludkovski (2005) for the computation) the copper mine example analyzed in Brennan and Schwartz (1985). Moreover, we can easily add further constraints to their model.

A related application is production of oil from oilfields. In the latter context,  $C_t$  can be increased by further exploration; at the same time fixed extraction costs increase as the field gets depleted, so that  $K_i \equiv K_i(C_t)$ . One may also add a termination option that allows total shutdown and avoids the ongoing O&M costs  $K_i$ .

**7.1.2. Hydroelectric Power Generation** This setting is similar to the pumped storage problem; however no pumping is available and the dammed reservoir is replenished solely with river run-off. The latter is modelled by a stochastic process  $(L_t)$ . When the turbines are running the produced electricity is sold at the spot power price  $G_t$ . As before, the inventory  $C_t$  is the current amount of water in the dammed reservoir. Reservoir management has been already studied by ancient Egyptians and Mesopotamians; related stochastic control models have recently been considered by Keppo (2002) and McNickle et al. (2004). Note that on a practical level a major challenge is the long-memory hydrological features of  $(L_t)$ .

**7.1.3. Power Supply Guarantees** Yet another possibility similar to the pumped storage above is the case of *power supply guarantees*. The latter involve a hybrid energy storage/power generation setting. By law, the North American Load-Serving Entities (LSE), i.e. the local power utilities, are obligated to provide power irrespective of demand. The latter is stochastic so that the LSE faces uncertain demand (volume risk), coupled with uncertain fuel prices (price risk). To insulate against risk, the LSE might operate an energy storage facility (e.g. a natural gas aquifer)

that acts as a buffer between risky supply costs and risky demand needs. Letting  $(D_t)$  represent the demand at time  $t$ , one obtains a model similar to (23) where the reservoir is depleted at rate  $D_t$  due to the requirement of producing fuel. Note that typically  $(D_t)$  is highly correlated with the fuel price  $(G_t)$  as high demand drives up the spot prices. Thus, the marginal cost of storage is high precisely when prices are high. See Deng et al. (2006) for further details.

**7.1.4. Emissions Trading** Another application area is emissions trading, cf. Insley (2003). A firm running a factory is subject to emission laws and must account for its pollution by buying publicly traded emission permits with current price  $G_t$ . The non-increasing inventory  $C_t$  in this case corresponds to the total number of remaining factory orders that must be filled in the current quarter. Hence, the management must satisfy all the orders while minimizing emission costs. The firm opportunistically runs its production given emission price  $G_t$  and current shipment backlog. This setup is similar to supply guarantees, with an additional constraint of  $C_t \leq C_0 - O_t$  where  $O_t$  is the (deterministic) shipping timetable supplied by the customers. Violation of this constraint causes a severe penalty as the firm misses its shipment.

Many other situations can be imagined—forest management, oilfield development, pipeline shipping, etc. From the above descriptions it should be evident that our numerical algorithm would carry over easily to the new settings. Summarizing, optimal switching with inventory is a widespread financial setting with many practical applications.

## 7.2. Incorporating Other Features.

From a practical standpoint the presented models are gross simplifications. However, as already advertised, the simulation framework permits great flexibility. To illustrate the possibilities, we briefly discuss additional features that a practitioner is likely to implement. First of all, one is likely to use a more general price model for  $(G_t)$  than (2). As already mentioned, all that is absolutely necessary is to satisfy Assumption 1, thus extra features such as regime-switching or latent factors are easily implementable. As already shown in Examples 2 and 3, seasonality effects, models with jumps and multi-factor models can be implemented straightforwardly.

The dynamics of  $(G_t)$  might also be affected by the choice of strategy. Indeed, since the manager tends to buy when prices are low and sell when prices are high, her influence is to smooth out the price fluctuations of  $(G_t)$ . This effect can be quite pronounced in segmented markets based on supply and demand (e.g. gas markets with little outside connectivity). As long as the effect is limited to the coefficients of (2),  $\mu = \mu(u_t, \cdot), \sigma = \sigma(u_t, \cdot)$ , such market impact can be treated by independently simulating price paths under each separate regime and then modifying the LSM

algorithm like in Carmona and Ludkovski (2008). If the transmission losses and engineering costs  $K_i$  are nonlinear in the pumping rates, it may become optimal to inject/withdraw gas at sub-maximal rates. In such a case, the optimal control  $u^*$  will take on a continuum of values. Our method relies heavily on  $u^*$  belonging to a (small) finite number of regimes; however a reasonable first-order correction would be to add a few more regimes (i.e. to discretize the range of  $u^*$ ) to the original three considered here.

Another challenge is proper modeling of borrowing constraints faced by the manager. In the North American gas industry, the facility typically borrows money in the summer to inject gas and then repays its loans in the winter as gas is withdrawn. In the meantime, the creditors often impose margin requirements regarding the value of stored gas versus the original loan. Thus, if prices drop, the manager might receive a margin call that would require him to sell off some of the inventory (at a loss) to raise capital. To account for this, one can let  $B_t$  be the cumulative borrowed capital taken out for inventory purchase. The margin constraints are then imposed on the net equity  $B_t - C_t \cdot G_t$ ; alternatively some absolute borrowing limits  $B_t > -K$  could be required.

The option of forward sales is another crucial feature. Forward sales allow the manager to lock-in future profits while reducing earnings volatility and form a bread-and-butter business in the gas storage industry. From the modeling perspective, a forward contract is challenging due to its non-Markovian nature, which necessitates complex account-keeping for gas already sold but still sitting in the facility (or gas already bought but still not on inventory). Indeed, imagine that at time  $t = t_0$ , the manager forward-sells some quantity  $C_0$  of gas at future date  $t = t_1$ . This now affects her possible future strategies: the manager must have at least  $C_0$  in inventory at  $t = t_1$ , and must also start to withdraw  $C_0$  after  $t_1$ . Such constraints are modelled easily enough in our framework; however because they affect the future, they are not easily implementable in the dynamic programming method —to find the value of the forward sale at  $t = t_0$  we must recompute the optimal strategy under the new admissibility restrictions  $\bar{U}(t, c, i)$ , which is computationally challenging (essentially requiring as much effort as the original computation). Hence, modeling of forward sales would lead to a “tree” of simulations with a separate branch for each possible forward sale or purchase. This might still be practical to do if the forward sales occur infrequently.

Finally, an interesting research direction would be incorporation of realistic risk objectives for the manager. In this paper we have assumed that the manager maximizes total (discounted) earnings from the asset. In practice this would lead to overly aggressive strategies and high earnings volatility. Thus, it is desirable to impose risk constraints that lead to more conservative speculation. One method for doing so can be achieved by replacing the linear conditional expectations in (12) with

non-linear expectations that take into account risk preferences Musiela and Zariphopoulou (2004). Alternatively, one may penalize the variance in cumulative gains/losses which were denoted by  $B_T$  above.

## 8. Conclusion.

This paper presented a simple model for energy storage that emphasizes the intertemporal optionality of the asset. Assuming that the commodity is bought and sold on the spot market, we have maximized the expected profit given operational constraints, in particular inventory limits and switching costs. While the model sidesteps the possibilities of forward trades, it properly accounts for the dynamic nature of the problem, which is a crucial aspect of revenue maximization.

Our approach is scalable and robust and we provide a detailed description of implementation. As our numerical examples attest, the model is computationally efficient and we believe better than any other proposed in the literature. We hope it can fill in the gap between current practitioner needs and academic models. Moreover, our strategy is applicable to many related problems, such as hydroelectric pumped storage, power supply guarantees, natural resource management and emissions trading.

**As the next step in improving our model, one should study more advanced risk objectives. In our model we have assumed linear risk-preferences of the agent who maximizes expected revenue under a given pricing measure. A real-life storage manager is in fact risk-averse, and would therefore take additional steps to hedge her risks (by e.g. reducing the variance of cashflows). However, direct hedging is likely to be of limited use in practice, because the facility buys gas based on *local* prices that are different from the liquidly traded indices. For instance, for gas storage the liquid instruments are the Henry Hub contracts available on the New York Mercantile Exchange (NYMEX). Such NYMEX trading would expose the agent to basis risk due to the local-index price spread. Consequently, to study hedging one must consider the risk-preferences of the manager, an issue that was alluded to in Section 7.2. On the theoretical level, financial hedging would require analysis of a combined control problem, namely the mixture of optimal switching and portfolio optimization in an incomplete market.**

## References

Ahn, H., A. Danilova, G. Swindle. 2002. Storing arb. *Wilmott* 1.

- ASCE. 1996. *Hydroelectric pumped storage technology : international experience*. New York, NY. Prepared by Task Committee on Pumped Storage of the Committee on Hydropower of the Energy Division of the American Society of Civil Engineers.
- Barrera-Esteve, C., F. Bergeret, C. Dossal, E. Gobet, A. Meziou, R. Munos, D. Reboul-Salze. 2006. Numerical methods for the pricing of Swing options: a stochastic control approach. *Methodology and Computing in Applied Probability* **8**(4) 517–540.
- Bensoussan, A., J.-L. Lions. 1984. *Impulse Control and Quasi-Variational Inequalities*. Gauthier-Villars, Paris.
- Blöchliger, L., T. Bollinger, J. Maurer, M. Semadeni. 2004. An evaluation model for the marked-to-market value of hydropower plants. *Proceedings of the 6th IAAE European Conference on Modelling in Energy Economics and Policy*.
- Brennan, M., E. Schwartz. 1985. Evaluating natural resource investments. *Journal of Business* **58** 135–157.
- Cao, M., A. Li, J. Wei. 2004. Precipitation modeling and contract valuation: A frontier in weather derivatives. *Journal of Alternative Investments* **7**(2) 92–99.
- Carmona, R., M. Ludkovski. 2008. Pricing asset scheduling flexibility using optimal switching. *Applied Mathematical Finance* **in press**. URL <http://www.informaworld.com/10.1080/13504860802170507>.
- Cartea, Alvaro, Thomas Williams. 2008. Uk gas markets: The market price of risk and applications to multiple interruptible supply contracts. *Energy Economics* **30**(3) 829–846.
- Chen, Z., P. A. Forsyth. 2007. A semi-lagrangian approach for natural gas storage valuation and optimal operations. *SIAM Journal on Scientific Computing* **30**(1) 339–368.
- Dayanik, S., M. Egami. 2005. A direct solution method for optimal switching problems of one-dimensional diffusions. Working paper.
- de Jong, C., K. Walet. 2003. To store or not to store. *Energy Power Risk Management* 8–11.
- Deng, S., Y. Oum, S.S. Oren. 2006. Hedging quantity risks with standard power options in a competitive wholesale electricity market. *Naval Research Logistics* **53**(7) 697–712.
- Doege, J., H.-J. Lüthi, Ph. Schiltknecht. 2006. Risk management of power portfolios and valuation of flexibility. *OR Spectrum* **28**(2) 267–287.
- Egloff, D. 2005. Monte Carlo algorithms for optimal stopping and statistical learning. *Annals of Applied Probability* **15**(2) 1396–1432.
- Eydeland, A., K. Wolyniec. 2003. *Energy and Power Risk Management: New Developments in Modeling, Pricing and Hedging*. John Wiley & Sons, Hoboken, NJ.
- FERC. 2004. State of and issues concerning underground natural gas storage. Tech. rep., Federal Energy Regulatory Commission. Packet no. AD04-11-0000.

- Geman, H. 2005. *Commodities and Commodity Derivatives : Modelling and Pricing for Agriculturals, Metals and Energy*. John Wiley& Sons, Hoboken, NJ.
- Glasserman, P., B. Yu. 2004. Number of paths versus number of basis functions in American option pricing. *Annals of Applied Probability* **14**(4) 2090–2119.
- Insley, M. 2003. On the option to invest in pollution control under a regime of tradable emissions allowances. *Canadian Journal of Economics* **35**(4) 860–883.
- Jaillet, P., E. Ronn, S. Tompaidis. 2004. Valuation of commodity based Swing options. *Management Science* **50**(7) 909–921.
- Keppo, J. 2002. Optimality with hydropower system. *IEEE Transactions on Power Systems*. **17**(3) 583–589.
- Kluge, T. 2004. Pricing options in electricity markets. Preprint.
- Longstaff, F.A., E.S. Schwartz. 2001. Valuing American options by simulations: a simple least squares approach. *Review of Financial Studies* **14** 113–148.
- Ludkovski, M. 2005. Optimal switching with application to energy tolling agreements. Ph.D. thesis, Princeton University.
- McNickle, D., E. Read, J. Tipping. 2004. The incorporation of hydro storage into a spot price model for the New Zealand electricity market. *Sixth European Energy Conference: Modelling in Energy Economics and Policy*. Zurich.
- Musiela, M., T. Zariphopoulou. 2004. A valuation algorithm for indifference prices in incomplete markets. *Finance & Stochastics* **8**(3) 399–414.
- Nowak, M., W. Römisich. 2000. Stochastic Lagrangian relaxation applied to power scheduling in a hydro-thermal system under uncertainty. *Annals of Operations Research* **100**(4) 251–272.
- Øksendal, B., A. Sulem. 2005. *Applied stochastic control of jump diffusions*. Springer-Verlag, Berlin.
- Pham, H., V. Ly Vath. 2007. Explicit solution to an optimal switching problem in the two-regime case. *SIAM J. Control Optim.* **46**(2) 395–426.
- Pirrong, C. 1996. *Corners and squeezes: the economics, law, and public policy of financial and commodity market manipulation*. Kluwer Academic Publishers.
- Porchet, A., N. Touzi, X. Warin. 2008. Valuation of a power plant under production constraints and market incompleteness. *Mathematical Methods of Operations research* (to appear).
- Stentoft, L. 2004. Assessing the least squares Monte Carlo approach to American option valuation. *Review of Derivatives Research* **7**(3) 129–168.
- Thompson, M., M. Davison, H. Rasmussen. 2003. Natural gas storage valuation and optimization: A real options approach. Tech. rep., University of Western Ontario.
- Tsitsiklis, J.N., B. van Roy. 2001. Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks* **12**(4) 694–703.

- 
- Wilmott, P., S. Howison, J. Dewynne. 1995. *The mathematics of financial derivatives*. Cambridge University Press, Cambridge. A student introduction.
- Zervos, M. 2003. A problem of sequential entry and exit decisions combined with discretionary stopping. *SIAM Journal on Control and Optimization* **42**(2) 397–421.